

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application Serial No. 10/821,832
Filing Date April 9, 2004
Inventorship Rajiv K. Grover
Applicant/Appellant Hewlett-Packard Company
Group Art Unit..... 2195
Examiner WILSER, Michael P.
Confirmation No. 1395
Attorney's Docket No. 200402482-1
Title: Device Loading In Storage Networks

APPEAL BRIEF

To: MS Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

As required under 37 C.F.R. §41.37(a), this brief is filed pursuant to the Office Action mailed April 29, 2009 reopening prosecution after Appeal, and is in furtherance to the Notice of Appeal filed herewith.

This brief contains items under the following headings as required by 37 C.F.R. §41.37 and M.P.E.P. §1206:

- I. Real Party In Interest
- II. Related Appeals, Interferences, and Judicial Proceedings
- III. Status of Claims
- IV. Status of Amendments
- V. Summary of Claimed Subject Matter
- VI. Grounds of Rejection to be Reviewed on Appeal
- VII. Argument
- VIII. Claims Appendix
- IX. Evidence Appendix
- X. Related Proceedings Appendix

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS, INTERFERENCES, AND JUDICIAL PROCEEDINGS

There are no other appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

A. Total Number of Claims in Application

There are twenty-six (26) claims pending in this application (Claims 1-26).

B. Current Status of Claims

1. Claims canceled: none.
2. Claims withdrawn from consideration but not canceled: none.
3. Claims pending: 1-26.
4. Claims allowed: none.
5. Claims rejected: 1-26.

C. Claims on Appeal

The claims on appeal are claims 1-26.

IV. STATUS OF AMENDMENTS

Appellant filed claim amendments in a Response after Final filed pursuant to the Examiner's request. Therefore the claims on appeal (as reflected in the claim appendix) are the claims previously presented in the Amendment after Final and are believed to have already been entered by the Examiner.

V. SUMMARY OF CLAIMED SUBJECT MATTER

According to claim 1, a computer program product (744, 748, and 752 in FIG. 7; p. 23, l. 1 to p. 25, l. 19) encoding a computer program (760 and 762 in FIG. 7; p. 23, l. 1 to p. 25, l. 19) for executing a computer process (FIG. 5 and FIG. 6; p. 20, l. 14 to p. 22, l. 11) on a computer system (730 in FIG. 7; p. 23, l. 1 to p. 25, l. 19). The computer process (FIG. 5 and FIG. 6) comprising identifying a plurality of storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4; p. 5, l. 6; p. 7, ll. 12-18; p. 8, ll. 3-8; p. 11, l. 7 to p. 12, l. 19; p. 16, ll. 1-16) to be configured in a storage network (100 in FIG. 1, and 400 in FIG. 4; p. 5, l. 13 to p. 7, l. 11; p. 15, ll. 18-21), identifying a number of host port Logical Unit Numbers (LUNs) (112a, 112b in FIG. 1 and 312, 317 in FIG. 3; p. 6, ll. 7-11; p. 7, ll. 1-6; p. 8, ll. 15-18; p. 11, ll. 12-16) which are configured on each of the storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4), identifying a number of host port connections to the storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4), and for each host port connection, determining actual loading of input/output (IO) jobs for each of the storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4)

based at least in part on a queue depth for each of the host port LUNs (112a, 112b in FIG. 1 and 312, 317 in FIG. 3).

According to claim 6, a computer program product (744, 748, and 752 in FIG. 7; p. 23, l. 1 to p. 25, l. 19) encoding a computer program (760 and 762 in FIG. 7; p. 23, l. 1 to p. 25, l. 19) for executing a computer process (FIG. 5 and FIG. 6; p. 20, l. 14 to p. 22, l. 11) on a computer system (730 in FIG. 7; p. 23, l. 1 to p. 25, l. 19). The computer process (FIG. 5 and FIG. 6) comprising identifying a plurality of storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4; p. 5, l. 4 to p. 6, l. 6; p. 7, ll. 12-18; p. 8, ll. 3-8; p. 11, l. 7 to p. 12, l. 19; p. 16, ll. 1-16) to be configured in a storage network (100 in FIG. 1, and 400 in FIG. 4; p. 5, l. 13 to p. 7, l. 11; p. 15, ll. 18-21), identifying a number of host port connections to the storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4), and for each host port connection, determining actual loading for each of the storage devices (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4) based at least in part on a queue depth for each of the host port connections.

According to claim 11, a method providing an input/output (IO) flow control mechanism in a storage network (100 in FIG. 1, and 400 in FIG. 4; p. 5, l. 13 to p. 7, l. 11; p. 15, ll. 18-21). The method comprising configuring a storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4; p. 5, l. 4 to p. 6, l. 6; p. 7, ll. 12-18; p. 8, ll. 3-8; p. 11, l. 7 to p. 12, l. 19; p. 16, ll. 1-16) in the storage network (100 in FIG. 1, and 400 in FIG. 4) with a plurality of host port Logical Unit Numbers (LUNs) LUNs (112a, 112b in FIG. 1 and 312, 317 in FIG. 3; p. 6, ll. 7-11; p. 7, ll. 1-6; p. 8, ll. 15-18; p. 11, ll. 12-16). The method also comprising identifying a queue depth for each of the host port LUNs (112a, 112b in FIG. 1 and 312, 317 in FIG. 3), automatically determining actual loading for the storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4) based at least in part on a queue depth for each host port LUN (112a, 112b in FIG. 1 and 312, 317 in FIG. 3). The

method also comprising accepting the storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4) configuration if the actual loading for the storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4) is no more than a maximum loading for the storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4).

According to claim 20, a method of device loading in a storage network (100 in FIG. 1, and 400 in FIG. 4; p. 5, l. 13 to p. 7, l. 11; p. 15, ll. 18-21). The method comprising configuring the storage network (100 in FIG. 1, and 400 in FIG. 4) with a plurality of host port connections to at least one storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4; p. 5, l. 14 to p. 6, l. 6; p. 7, ll. 12-18; p. 8, ll. 3-8; p. 11, l. 7 to p. 12, l. 19; p. 16, ll. 1-16). The method also comprising for each of a plurality of host port connections to the at least one storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4), determining actual loading of the at least one storage device (110a-c in FIG. 1, 210 in FIG. 2, 310 and 315 in FIG. 3, 450 in FIG. 4) based at least in part on a queue depth of each host port connection so that the number of input/output (IO) jobs being issued by a host do not exceed the queue depth of a service queue.

The summary is set forth in several exemplary embodiments that correspond to the independent claims. It is noted that no dependent claims containing means plus function are argued separately. Discussions about elements and recitations to these claims can be found at least at the cited locations in the specification and drawings.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The Office Action rejected claims 1-2, 4, 6-7, 9, 11-14, 16-21, 23-24, and 26 under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 7,225,242 to Cherian, et al. ("Cherian") in view of the IBM publication to Tate,

et al. ("Tate"). The Final Office Action also rejected claims 3, 5, 8, 10, 15, 22, and 25 under 35 U.S.C. 103(a) as being unpatentable over Cherian and Tate and further in view of U.S. Patent No. 2004/0078599 to Nahum ("Nahum"). Appellant requests the Board to review each of these grounds of rejection.

VII. ARGUMENT

First Rejection under 35 U.S.C. §103(a)

Claims 1-2, 4, 6-7, 9, 11-14, 16-21, 23-24, and 26 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Cherian in view of Tate.

In its decision, *KSR Int'l Co. v. Teleflex, Inc.*, No 04-1350 (U.S. Apr. 30, 2007), the Supreme Court reaffirmed application of the Graham factors in making a determination of obviousness under 35 U.S.C. § 103(a). The four factual inquiries under Graham are: (1) determining the scope and contents of the prior art; (2) ascertaining the differences between the prior art and the claims in issue; (3) resolving the level of ordinary skill in the pertinent art; and (4) evaluating evidence of secondary consideration. Even if all of the prior art elements are disclosed by separate prior art references, the Examiner still must identify the reason why a person of ordinary skill in the art would have combined the prior art elements in the manner claimed.

Independent Claims 1 and 6

Claim 1 recites "for each host port connection, determining actual loading of IO jobs for each of the storage devices based at least in part on a queue depth for each of the host port LUNs" (emphasis added).

The Examiner relies on Cherian as teaching these recitations. Cherian discloses correlating the command queue depth of the servers of a SAN to the

command queue depth of the storage controllers of the computer network. For each storage controller, the execution throttle levels for those servers that have logical ownership over logical storage units of the storage controller are verified. If the summed execution throttle level exceeds the command queue depth, the execution throttle level of one or more of the servers of the network is adjusted to insure that the potential command throughput of the servers conform to the standards established for governing the relation of command throughput between servers and storage controllers. *See Abstract.*

Execution throttle levels and potential command throughput of the servers is not the same as the actual loading for each of the storage devices.

Claim 6 is rejected for the same reasons as claim 1 and therefore is believed to be allowable for at least the same reasons as claim 1.

For at least the foregoing reasons, the Examiner has failed to establish that independent claim 1 is obvious in view of Hinton and Tawil.

Dependent Claim 2 and 7

Claim 2 depends from claim 1, which is believed to be allowable. Therefore, claim 2 is also believed to be allowable for at least the same reasons as claim 1.

In addition, claim 2 recites "determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network." The Examiner relies on col. 5, lines 23-27 in Cherian as teaching these recitations. Here, Cherian discloses:

(Equation 1) Execution Throttle on Servers $(A+B+C) < \text{Command Queue Depth of Storage Device } X$.

(Equation 2) Execution Throttle on Servers $(B+C+D+E) < \text{Command Queue Depth of Storage Device Y}$

Again, however, the execution throttle for various servers in Cherian is not the same as determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network.

Claim 7 depends from claim 6, which is believed to be allowable. Therefore, claim 7 is also believed to be allowable for at least the same reasons as claim 6. Claim 7 is also believed to be allowable for the additional reasons set forth above for claim 2.

Dependent Claims 4 and 9

Claim 4 depends from claim 1, which is believed to be allowable. Therefore, claim 4 is also believed to be allowable for at least the same reasons as claim 1.

Claim 4 recites “uses a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading.” The Examiner relies on col. 5, lines 21-44 in Cherian. In addition to Equations 1 and 2 discussed above, Cherian also discloses:

As to Equation 1, because each of Servers A, B, and C has logical ownership of a LUN on Storage Device X, the execution throttle of each of Servers A, B, and C is summed and compared to the command queue depth of Storage Device X. As to Equation 2, because each of Servers B, C, D, and E has logical ownership of a LUN on Storage Device Y, the execution throttle levels of each of

Servers B, C, D, and E are summed and compared to the command queue depth of Storage Device Y.

The verification step of step 22 is performed for each storage controller of the SAN. If the verification step is not satisfied for any storage controller, processing continues at step 26, where the execution throttle is incremented or decremented by at least one increment on at least one of the servers associated with the storage controller associated with the unsatisfied verification step.

Summing the execution throttle for various servers in Cherian is not the same as using a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading.

Claim 9 depends from claim 6, which is believed to be allowable. Therefore, claim 9 is also believed to be allowable for at least the same reasons as claim 6. Claim 9 is also believed to be allowable for the additional reasons set forth above for claim 4.

Independent Claim 11

Claim 11 is rejected for the same reasons as claim 1 and therefore is believed to be allowable for at least the same reasons as claim 1.

In addition, claim 11 recites "A method providing an input/output (IO) flow control mechanism in a storage network." The Examiner provided no support for rejecting this recitation.

If the Examiner wants to dismiss these recitations as being part of the preamble, it is error for the Examiner to summarily dismiss these recitations simply on the basis that these recitations are a part of the preamble. The MPEP

states at Section 2111.02 (Effect of Preamble) that “[t]he determination of whether a preamble limits a claim is made on a case-by-case basis in light of the facts in each case; there is no litmus test defining when a preamble limits the scope of a claim.”

The MPEP goes on in Section 2111.02 to state that “[t]he claim preamble must be read in the context of the entire claim. The determination of whether preamble recitations are structural limitations or mere statements of purpose or use “can be resolved only on review of the entirety of the [record] to gain an understanding of what the inventors actually invented and intended to encompass by the claim.”

During examination, statements in the preamble reciting the purpose or intended use of the claimed invention must be evaluated to determine whether the recited purpose or intended use results in a structural difference between the claimed invention and the prior art. In this case, the recitations in claim 11 of “providing an input/output (IO) flow control mechanism in a storage network” clearly is a structural difference between the claimed invention and the prior art. Therefore, it was error for the Examiner to dismiss these recitations during examination.

For at least the foregoing reasons, the Examiner has failed to establish that independent claim 11 is obvious in view of Hinton and Tawil.

Dependent Claims 12-14 and 16-17

Claims 12-14, and 16-17 depend from claim 11, which is believed to be allowable. Therefore, claims 12-14, and 16-17 are also believed to be allowable for at least the same reasons as claim 11.

In addition, claims 12-14 and 16-17 include further recitations for “automatically determining actual loading for the storage device” which is not taught by the cited references, as already discussed above.

Dependent Claim 18

Claim 18 depends from claim 11, which is believed to be allowable. Therefore, claim 18 is also believed to be allowable for at least the same reasons as claim 11.

In addition, claim 18 recites “wherein the maximum loading for the storage device is based on a loading factor.” The Examiner summarily dismisses these recitations as being obvious “to ensure that the system will perform with minimum error.” Claim 19 recites “the loading factor is in the range of about 80% to 90% of the service queue depth for the storage device.” The Examiner summarily dismisses these recitations as being obvious “that the range should not be 100% of the service queue depth, and should not be too low, thus, 80-90% would be a safe choice for the system.” There is no support in the references for either of these conclusions. Clearly this is nothing more than hindsight interpretation of the references in view of the claim recitations.

Dependent Claim 19

Claim 19 depends from claim 11, which is believed to be allowable. Therefore, claim 19 is also believed to be allowable for at least the same reasons as claim 11.

Independent Claim 20

Claim 20 is rejected for the same reasons as claim 1 and therefore is believed to be allowable for at least the same reasons as claim 1.

Dependent Claim 21

Claims 21 depends from claim 20, which is believed to be allowable. Therefore, claim 21 is also believed to be allowable for at least the same reasons as claim 20.

In addition, claim 21 is rejected for the same reasons as claim 2. Therefore, claim 21 is also believed to be allowable for the further reasons discussed above for claim 2.

Dependent Claims 23-24

Claims 23-24 depend from claim 20, which is believed to be allowable. Therefore, claims 23-24 are also believed to be allowable for at least the same reasons as claim 20.

In addition, claims 23-24 are rejected for the same reasons as claims 16-17. Therefore, claims 23-24 are also believed to be allowable for the further reasons discussed above for claims 16-17.

Dependent Claim 26

Claim 26 depends from claim 20, which is believed to be allowable. Therefore, claim 26 is also believed to be allowable for at least the same reasons as claim 20.

Second Rejection under 35 U.S.C. §103(a)

Claims 3, 5, 8, 10, 15, 22, and 25 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Cherian and Tate and further in view of Nahum.

In its decision, *KSR Int'l Co. v. Teleflex, Inc.*, No 04-1350 (U.S. Apr. 30, 2007), the Supreme Court reaffirmed application of the Graham factors in making a determination of obviousness under 35 U.S.C. § 103(a). The four factual inquiries under Graham are: (1) determining the scope and contents of the prior art; (2) ascertaining the differences between the prior art and the claims in issue; (3) resolving the level of ordinary skill in the pertinent art; and (4) evaluating evidence of secondary consideration. Even if all of the prior art elements are disclosed by separate prior art references, the Examiner still must identify the reason why a person of ordinary skill in the art would have combined the prior art elements in the manner claimed.

Dependent Claim 3

Claim 3 depends from claim 1, which is believed to be allowable. Therefore, claim 3 is also believed to be allowable for at least the same reasons as claim 1.

In addition, claim 3 recites “determining actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network.” The Examiner admits that Cherian and Tate do not disclose these recitations. Applicant agrees with this admission. However, the Examiner relies on paragraph [0018] in Nahum as disclosing these recitations. Appellant disagrees.

Paragraph [0018] describes a security procedure for authenticating each host. Nahum does not determine actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network.

Dependent Claims 5 and 10

Claim 5 depends from claim 1, which is believed to be allowable. Claim 10 depends from claim 6, which is believed to be allowable. Therefore, claims 5 and 10 are also believed to be allowable for at least the same reasons as claims 1 and 6, respectively.

In addition, claim 5 recites “the computer process further simplifies host groups and LUN security groups into virtual connections for analysis.” Claim 10 includes similar recitations as claim 5. Again the Examiner cites broadly to paragraph [0018] in Nahum as disclosing these recitations. However, Applicant cannot find any basis for the rejection in paragraph [0018]. The Examiner has yet to clarify this position.

The Examiner previously stated that “Applicant failed to provide any reason why the cited passages do not teach the claim limitations, and therefore are not persuasive.” The burden is initially on the Patent Office to make a prima facie case in support of any rejection. This burden does not shift to the Applicant until and unless the Patent Office has met its burden. Here, Appellant asserts that the broad reference to paragraphs [0018] failed to support the rejection as it is not specifically called out what in paragraph [0018] is being relied on as “the computer process further simplifies host groups and LUN security groups into virtual connections for analysis” in claim 5. Applicant also has reviewed paragraph [0018] and cannot understand what is being relied on as teaching or suggestion each of these claim recitations. The Examiner has refused to clarify this position.

Dependent Claims 8, 10, and 15

Claims 8 and 10 depend from claim 6, which is believed to be allowable. Claim 15 depends from claim 11, which is believed to be allowable. Therefore,

claims 8, 10, and 15 are also believed to be allowable for at least the same reasons as claim 6 and 11, respectively.

Dependent Claim 22

Claim 22 depends from claim 20, which is believed to be allowable. Therefore, claim 22 is also believed to be allowable for at least the same reasons as claim 20.

Dependent Claim 25

Claim 25 depends from claim 20, which is believed to be allowable. Therefore, claim 25 is also believed to be allowable for at least the same reasons as claim 20.

In addition, claim 25 was rejected on the same basis as claims 3 and 26. Therefore, claim 25 is also believed to be allowable for at least the same reasons already discussed above for these claims.

Conclusion

For the reasons provided herein, Appellant respectfully requests the Board to rule that the rejections of the claims are improper.

Respectfully Submitted,

/Mark D. Trenner/

Dated: July 29, 2009

By: _____

Mark D. Trenner

Reg. No. 43,961

(720) 221-3708

VIII. CLAIMS APPENDIX

1. A computer program product including computer-readable storage with a computer program, the computer program executing a computer process on a computer system, the computer process:

identifying a plurality of storage devices to be configured in a storage network;

identifying a number of host port Logical Unit Numbers (LUNs) which are configured on each of the storage devices;

identifying a number of host port connections to the storage devices; and

for each host port connection, determining actual loading of input/output (IO) jobs for each of the storage devices based at least in part on a queue depth for each of the host port LUNs.

2. The computer program product of claim 1 wherein the computer process further comprises determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network.

3. The computer program product of claim 1 wherein the computer process further comprises determining actual loading for each of the storage devices based at least in part on a number of LUN security groups in the storage network.

4. The computer program product of claim 1 wherein the computer process further uses a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading.

5. The computer program product of claim 1 wherein the computer process further simplifies host groups and LUN security groups into virtual connections for analysis.

6. A computer program product including computer-readable storage with a computer program, the computer program executing a computer process on a computer system, the computer process:

identifying a plurality of storage devices to be configured in a storage network;

identifying a number of host port connections to the storage devices; and

for each host port connection, determining actual loading for each of the storage devices based at least in part on a queue depth for each of the host port connections.

7. The computer program product of claim 6 wherein the computer process further comprises determining actual loading for each of the storage devices based at least in part on a number of host groups in the storage network.

8. The computer program product of claim 6 wherein the computer process further comprises determining actual loading for each of the storage devices based at least in part on a number of Logical Unit Numbers (LUNs) LUN security groups in the storage network.

9. The computer program product of claim 6 wherein the computer process further uses a loading factor to determine if the actual loading for each of the storage devices exceeds a maximum loading.

10. The computer program product of claim 6 wherein the computer process further simplifies host groups and LUN security groups into virtual connections for analysis.

11. A method providing an input/output (IO) flow control mechanism in a storage network, comprising:

configuring a storage device in the storage network with a plurality of host port Logical Unit Numbers (LUNs);

identifying a queue depth for each of the host port LUNs;
automatically determining actual loading for the storage device based at least in part on a queue depth for each host port LUN; and
accepting the storage device configuration if the actual loading for the storage device is no more than a maximum loading for the storage device.

12. The method of claim 11 wherein automatically determining actual loading for the storage device is also based at least in part on a number of host paths connected to the storage device.

13. The method of claim 11 wherein automatically determining actual loading for the storage device port is also based at least in part on a number of LUNs configured for the storage device.

14. The method of claim 11 wherein automatically determining actual loading for the storage device is also based at least in part on a number of host groups in the storage network.

15. The method of claim 11 wherein automatically determining actual loading for the storage device is also based at least in part on a number of LUN security groups in the storage network.

16. The method of claim 11 further comprising automatically determining actual loading for a plurality of backend LUNs connected to the storage device.

17. The method of claim 11 further comprising iteratively determining actual loading for a plurality of storage devices in the storage network.

18. The method of claim 11 wherein the maximum loading for the storage device is based on a loading factor for test environments.

19. The method of claim 18 wherein the loading factor is in the range of about 80% to 90% of a service queue depth for the storage device.

20. A method of device loading in a storage network, comprising:

 configuring the storage network with a plurality of host port connections to at least one storage device; and

 for each of a plurality of host port connections to the at least one storage device, determining actual loading of the at least one storage device based at least in part on a queue depth of each host port connection so that the number of input/output (IO) jobs being issued by a host do not exceed the queue depth of a service queue.

21. The method of claim 20 wherein determining actual loading is also based at least in part on a number of host groups in the storage network.

22. The method of claim 20 wherein determining actual loading is also based at least in part on a number of Logical Unit Numbers (LUNs) LUN security groups in the storage network.

23. The method of claim 20 further comprising determining actual loading for a plurality of backend LUNs connected to the at least one storage device.

24. The method of claim 20 further comprising iteratively determining actual loading for a plurality of storage devices in the storage network.

25. The method of claim 20 wherein device loading is a function of maximum queue depth for each target port, number of host paths connected to the target port, queue depth for each LUN on the host port, and number of LUNs configured on the target port.

26. The method of claim 20 wherein device loading is a function of queue depth for each target port, number of host paths connected to the target port, and queue depth for each host port.

IX. EVIDENCE APPENDIX

Not applicable.

X. RELATED PROCEEDINGS APPENDIX

Not applicable.